

ABSTRACT

Techniques for recovering an original image from its halftone version using information obtained from the dither matrix used to generate the halftone. The techniques involve determining a bounding region defined by (x_{\min}, x_{\max}) for each pixel location and color value, based on the information obtained from the dither matrix, and applying a low-pass filter to the halftone version as follows. For each pixel color, the post-low-pass-filtered value of that color is used, if the filtered value is greater than or equal to x_{\min} and less than or equal to x_{\max} , the x_{\min} value is used, if the filtered value is less than x_{\min} , and the x_{\max} value is used, if the filtered value is greater than x_{\max} . Next, it is determined, for each pixel, whether its value for that color is to be used in an averaging process based on the distance between the bounding region of that pixel and neighborhood pixels' bounding regions.